

Our Ref.: 002717.P046

APPLICATION FOR UNITED STATES LETTERS PATENT

FOR

**Method and System for Increasing Participation
in a Standby Router Protocol**

Inventor(s): **Michael Yip**
Apoorva Bhatt

Prepared by:

BLAKELY SOKOLOFF TAYLOR & ZAFMAN, LLP
12400 Wilshire Boulevard, 7th Floor
Los Angeles, California 90025
(503) 684-6200

"Express Mail" label number EL034436965US

**Method and System for Increasing Participation
in a Standby Router Protocol**

5 **BACKGROUND OF THE INVENTION**

1. Field of the Invention

The present invention relates to the field of network management technologies. In particular, the present invention relates to increasing the numbers of end-hosts that can participate in a standby router protocol.

10 2. Background Information and Description of Related Art

The use of standby routers in an Internet Protocol (IP) network is known in the art. The Internet Engineering Task Force (IETF) has published a draft standard protocol for using standby routers, also referred to as redundant routers, entitled *Virtual Router Redundancy Protocol*, version 2-05, on January 5, 2000 (VRRP).

15 In a typical network configuration, end-hosts that are connected to a layer-2 domain communicate with other subnets through the use of a default router. Often, the default router is statically configured as it minimizes configuration and processing overhead on the end-host and is widely supported by most Internet Protocol (IP) networks. As noted by the IETF, one of the drawbacks of using a
20 statically configured default router is that it creates a single point of failure.

Therefore, loss of the default router results in a catastrophic event, isolating all end-hosts that are unable to detect any alternate path that may be available. The use of standby routers, also referred to as redundant routers, eliminates the single point of

failure inherent in the static default routed environment. (VRRP, Section 1, Introduction).

Protocols for using standby routers involve the notion of a virtual router. A virtual router is an abstract object managed by a standby router protocol (SRP), and it functions as a default router for end-hosts on a network. The virtual router is defined by a Virtual Router Identifier (VRID) and a set of associated IP addresses. The virtual router may be implemented with two or more routers running the SRP. The SRP specifies an election process whereby the responsibility for forwarding packets sent to the IP address(es) associated with the virtual router is dynamically assigned to one of the SRP routers, called the master. The remaining SRP routers are variously referred to as standby, backup, or slave routers, and are available to assume forwarding responsibility for a virtual router should the current master fail.

Figure 1A is a block diagram that illustrates a typical prior art network configuration using the IETF VRRP. As illustrated, routers R1 110 and R2 120 are defined as VRRP routers connected to a virtual local area network VLAN1 115 supporting virtual routers VRID1 and VRID2. VRID1 is defined as the virtual router associated with IP subnet 10.2.3.1, and VRD2 is defined as the virtual router associated with IP subnet 10.2.4.1. Hosts H1 130 and H2 135 have configured a static default route through R1's IP address 10.2.3.1, and hosts H3 140 and H4 145 have configured a static default route through R2's IP address 10.2.4.1. R1 110 is the initial master for VRID1 and R2 120 is the backup (slave) router. Likewise, R2 is the initial master for VRID2 and R1 is the backup (slave) router. Thus, if R1 110 fails such as when the R1 ISP 150 connection 111 to the Internet 155 goes down, then R2 120 is elected the new master VRRP router for VRID1, and publishes the new subnet route 10.2.4.1 for hosts H1 130 and H2 135. Likewise, if R2 120 fails

Yip et al. – Method and System for Increasing Participation
in a Standby Router Protocol

EL034436965US

DJC/mwb

(e.g. the connection 121 to the Internet 155 goes down), then R1 110 is elected the new master VRRP router for VRID2, and publishes the new subnet route 10.2.3.1 for hosts H3 140 and H4 145. The election of the new master VRRP router is performed in accordance with the election process defined for the IETF VRRP
5 protocol.

One of the drawbacks to implementing an SRP is that the SRP messaging that is necessary to support the election process generates a significant amount of network traffic. SRP messaging is performed using Internet Protocol (IP) multicast datagrams, specifically referred to as SRP packet datagram units (PDUs). Each
10 end-host, subnet or any layer-2 domain participating in the SRP must send a PDU containing information about their status to the two or more routers running the SRP. If a large number of end-hosts, subnets, or layer-2 domains participate, the result is a periodic flooding of the network with SRP PDUs to and from the SRP routers.

Figure 1B is a block diagram that illustrates a typical prior art network configuration similar to that illustrated in **Figure 1A**, but with 200 VLANs 160, 161, 162, and 163 participating in the VRRP instead of a single VLAN1 115. Because the router R1 110 and R2 120 must exchange PDUs for all 200 VLANs over connection 115, the network can be overwhelmed. Nonetheless, all 200 VLANs
15 must participate in the SRP in order to provide the benefits of dynamic failover that is critical to the support of a high level of uninterrupted service on a network. In some cases, it may be necessary to support thousands of VLANs or other layer-2 domains. It would be desirable, therefore, to devise a way to increase participation in an SRP without a concomitant increase in the amount of network traffic due to
20 SRP messaging in order to provide scalability when implementing the SRP.
25

Another problem when implementing an SRP is that host-specific ports on the SRP routers running the SRP (i.e. the routers comprising the virtual router), are not utilized in a typical SRP routing configuration. Host-specific ports are generally used by a single end-host for which the port is specifically configured.

5 **Figure 1C** is a block diagram that illustrates a typical prior art network configuration similar to that illustrated in **Figure 1B**. In addition to the ports that support the 200 VLANs 160, 161, 162, and 163 that participate in the VRRP, routers R1 110 and R2 120 include host-specific ports that support end-host H5 112 and end-host H6 122, respectively. Because routers R1 and R2 are the end-hosts' only
10 connection to the rest of the network and the Internet, R1 and R2 are still a single point of failure for those end-hosts. As a result, when used at all, these host-specific ports end up being relegated to less critical end-hosts where the need for redundancy is not as imperative. It would be desirable, therefore, to devise a way to allow these under-utilized ports to participate in SRP.

SUMMARY

According to one aspect of the invention, a method and system is provided for increasing participation in a standby router protocol (SRP) without increasing the
20 amount of network traffic due to SRP messaging. One or more domain master VLANs participate in an SRP on behalf of numerous member VLANs belonging to the domain master VLAN's domain. The domain master VLANs are associated with at least one virtual router. Each domain member VLAN follows the current default routing configuration for the domain master VLAN as determined in accordance with
25 the SRP election process defined for that virtual router.

According to another aspect of the invention, an end-host attached to a host-specific port on an SRP router participates in the SRP by using a router bridge connection established between the SRP router and other SRP routers supporting the virtual router. Each end-host on a host-specific port follows the current default
5 routing configuration implemented by the current master SRP router as determined in accordance with the SRP election process defined for that virtual router. In some instances, the current master router may be the same SRP router having a host-specific port to which the end-host is attached, and in other cases, the current master router may be a different router from the one having a host-specific port to
10 which the end-host is attached.

In accordance with other aspects of the present invention, apparatus are provided for carrying out the above and other methods.

BRIEF DESCRIPTION OF DRAWINGS

15 The present invention will be described by way of exemplary embodiments, but not limitations, illustrated in the accompanying drawings in which like references denote similar elements, and in which:

Figure 1A illustrates a block diagram of a typical prior art network configuration using the virtual redundant router protocol (VRRP);

20 **Figure 1B** illustrates a block diagram of prior art network of **Figure 1A** with 200 VLANs participating in the VRRP;

Figure 1C illustrates a block diagram of prior art network of **Figure 1B** with end-hosts connected to host-specific ports that are unable to participate in the VRRP;

Figure 2 illustrates a block diagram of a network configuration using a standby router protocol (SRP) in accordance with one embodiment of the present invention;

Figure 3 illustrates a flow diagram of the participation in an SRP of a domain master VLAN in accordance with one embodiment of the present invention; and

Figure 4 illustrates a flow diagram of the participation in an SRP of an end-host connected to a host-specific port of an SRP router in accordance with one embodiment of the present invention.

10 DETAILED DESCRIPTION OF THE INVENTION

In the following description various aspects of the present invention, a method for increasing participation in an SRP will be described. Specific details will be set forth in order to provide a thorough understanding of the present invention. However, it will be apparent to those skilled in the art that the present invention may be practiced with only some or all of the described aspects of the present invention, and with or without some or all of the specific details. In some instances, well known architectures, steps, and techniques have not been shown to avoid unnecessarily obscuring the present invention. For example, specific details are not provided as to whether the method and system is implemented in a router, server or gateway, as a software routine, hardware circuit, firmware, or a combination thereof.

Various operations will be described as multiple discrete steps performed in turn in a manner that is most helpful in understanding the present invention. However, the order of description should not be construed as to imply that these operations are necessarily performed in the order they are presented, or even order

dependent. Lastly, repeated usage of the phrase “in one embodiment” does not necessarily refer to the same embodiment, although it may.

Referring now to **Figure 2**, wherein a block diagram of a network configuration 200 using a standby router protocol (SRP) in accordance with one embodiment of the present invention is shown. As illustrated, routers R3 210 and R4 220 are defined as SRP routers connected to a local area network (LAN) 114 supporting virtual routers VRID1 and VRID2. As in **Figures 1A-1C**, VRID1 is defined as the virtual router associated with IP subnet 10.2.3, and VRD2 is defined as the virtual router associated with IP subnet 10.2.4. Each IP subnet may be configured as one or more virtual LAN (VLAN) domains such as Domain Master VLAN A 230 and Domain Master VLAN B 240. Layer-2 subnets such as VLANs 1-200 160/161/162/163 and VLANs 201-400 164/165 are configured as members of one of the VLAN domains and follow the default routing rules currently in place for their Domain Master VLAN. In one embodiment, Domain Master VLAN A 230 has configured a static default route through R3's IP address 10.2.3.1 and member VLANs 1-200 160/161/162/163 follow the default routing rules for Domain Master VLAN A 230. Similarly, Domain Master VLAN B 240, which has configured a static default route through R4's IP address 10.2.4.1 and member VLANs 201-400 164/165 follow the default routing rules for Domain Master VLAN B 240.

End-hosts H5 112 and H6 122 are each connected to a Host-Specific Port 212/222 on their respective SRP routers R3 210 and R4 220. End-host H5, having an IP address of 10.2.3.9, has configured a static default route through R3's IP address 10.2.3.1, and end-host H6, having an IP address of 10.2.4.51, has configured a static default route through R3's IP address 10.2.4.1. A bridge route 226 is defined between R3 210 and R4 220 to provide a redundant path via external

connections 211/212 to the ISP 150 and the Internet 155. The current SRP master for a given virtual router, VRID 1 or VRID 2, serves as the default router for both of the end-hosts H5 112 and H6 122. In this way, if connection 211 on R3 210 goes down and R4 220 becomes the SRP master router, then end-host H5 may connect
5 to ISP 150 via the bridge route 226 to SRP router R4 220 and connection 221.

Similarly, if connection 221 on R4 220 goes down and R3 210 becomes the SRP master router, then end-host H6 122 may connect to ISP 150 via the bridge route 226 to SRP router R3 210 and connection 211. Without such redundancy, end-hosts H5 and H6 could become isolated from the outside network services of the

10 ISP 150 and the Internet 155 should their static default router connection to ISP 150 go down.

Additional layer-2 subnets may be connected to Domain Masters VLAN C, D, etc. 240, with each Domain Master VLAN having configured a static default route through the SRP router R3, R4, or any other SRP router connected to network 114.

15 In the illustrated embodiment, R3 210 is the initial master for VRID1 and R4 220 is the standby (backup or slave) router. Likewise, R4 220 is the initial master for VRID2 and R3 210 is the standby (backup or slave) router. Other initial master configurations may be employed without departing from the spirit of the invention.

Each router R3 210 and R4 220, is configured to run an SRP protocol in
20 accordance with one embodiment of the invention. Although only two SRP routers are shown in the illustrated embodiment, additional SRP routers may be added to network 114 to increase the level of redundancy. In one embodiment each router is a switch that is physically connected to the same layer-2 domain and the rest of the network 114. Each switch is configured with the same subnet information of the

25 Domain Master VLANs they are sharing so that either switch is capable of

functioning as the default router for a given Domain Master VLAN 230/240 and their associated member VLANs 160-165. Initially, one switch will be elected as the master SRP router, and the other will be the standby (backup or slave) router. For example, with reference to the illustrated embodiment, R3 210 is the master SRP router for VRID1. As the master, R3 functions as the default router for the subnet 10.2.3 associated with hosts H1 and H2. It will handle all the data traffic for hosts H1 and H2 and export the subnet route for H1 and H2 to other routers. Likewise, R4 220 is the master SRP router for VRID2. Thus, both VRID1 and VRID2 are each served by a master and a standby (backup or slave) router.

In operation, R3 210 and R4 220 periodically trade SRP packet data units (PDUs) 225 with each other at specified time intervals. The SRP PDUs contain, among other items, data that reflect the SRP election algorithm parameters obtained from participating Domain Master VLANs A, B, C, D etc. 230/240 for the routers R3 and R4. The election algorithm parameter data is used by the election process to determine the new master router. One example of the type of data and election algorithms used to determine the new master is described in commonly assigned application serial number 09/797,475, Method and Apparatus for Selecting Redundant Routers using Tracking. However, it should be noted that numerous other data and election algorithms may be employed without departing from the spirit of the invention.

If the SRP election algorithm parameter data for either of the SRP routers R3 210 and R4 220 indicate that a new master must be elected, than an election process elects the new master in accordance with the election algorithm for the SRP. For example, in one embodiment, R3 210 may be elected to assume the master router function for virtual router VRID2 serving subnet 10.2.4 and associated

Domain Master VLANs A, B, C, D etc. An election process is performed, and upon election of R3 210 to be the master router for virtual router VRID2, R3 210 sends the new subnet route for end-hosts of the VLAN members 160-165 of Domain Master VLANs A, B, C, D etc. 230/240 to other routers on the network, and
5 commences default routing. In this scenario, R3 210 will be the master for both virtual routers VRID1 and VRID2.

Referring now to **Figure 3**, wherein a flow diagram of the participation in an SRP of a domain master VLAN 300 in accordance with one embodiment of the present invention is shown. As illustrated, in block 310, the network administrator
10 establishes one or more Domain Master VLANs A and B (or C, D, etc.) as needed to support the end-hosts for which redundancy is desired. The R3 SRP router receives SRP PDUs from Domain Master VLAN A at block 320, and Domain Master VLAN B at block 330 and so forth until SRP PDUs are received from all participating Domain Master VLANs. At block 340, the R3 and R4 SRP routers, as well as any
15 other SRP routers that are supporting the various virtual routers, e.g. VRID1 and VRID2, periodically exchange the SRP PDUs obtained from the participating Domain Master VLANs A and B (or C, D, etc.). When the parameter data obtained from the SRP PDUs indicates that it is necessary to elect a new master for a given virtual router VRID1 or VRID2 or any other virtual router that the SRP routers are
20 supporting, an SRP election algorithm is performed to elect the new master SRP router accordingly.

For example, the parameter data may be in the form of priority field values which indicate the relative priority of one router over another for a given virtual router based on tracking environmental factors as described in an embodiment of
25 commonly assigned application Method and Apparatus for Selecting Redundant

Yip et al. – Method and System for Increasing Participation
in a Standby Router Protocol

Routers using Tracking, Application Serial No. 09/797,475. Other forms of parameter data may be obtained from the SRP PDUs without departing from the spirit of the invention.

At block 360, the newly elected master SRP router sends its subnet route for
5 Domain Master VLANs A and B (and C, D, etc.) to all of the other routers, and commences default routing for all Domain Master VLANs A and B (and C, D, etc.) and associated VLAN member subnets once the old master SRP router times out.

Referring now to **Figure 4**, wherein a flow diagram of the participation in an SRP of a Host Specific Port 400 in accordance with one embodiment of the present
10 invention is shown. As illustrated, in block 410, the network administrator establishes end-hosts on host-specific ports 212/222 on the R3 and R4 SRP routers respectively. For example, the network administrator may assign end-host H5 to host-specific port 212 on R3 420, or end-host H6 to host-specific port 222 on R4 430. The end-hosts H5 and H6 are now able to participate in the SRP for a given
15 virtual router, e.g. VRID1 and VRID2.

At block 440, the current SRP master router establishes a bridge route 226 from the host-specific port 212 on SRP router R3 210 to at least one other SRP router supporting the virtual routers VRID1 and VRID2, including router R4 220. Similarly, the current master router establishes a bridge route 226 from the host-
20 specific port 222 on SRP router R4 220 to at least one other SRP router supporting the virtual routers VRID1 and VRID2, including router R4 220. At block 450, the SRP master router sends the bridge route 226 to all other routers and commences default routing for end-hosts H5 and H6. Eventually, at block 460, the virtual router which the current SRP master router is supporting performs an election algorithm to
25 elect a new master SRP router. Block 440 is repeated to establish the new SRP

Yip et al. – Method and System for Increasing Participation
in a Standby Router Protocol

master router's bridge route for the host-specific ports on the SRP routers R3 and R4 (or other SRP routers support the virtual router). Block 450 is also repeated to send the re-established bridge route 226 to all other routers and the new SRP master router commences default routing for end-hosts H5 and H6 attached to the host-specific ports 212/222. The SRP continues in this fashion over time to provide redundancy and a resulting high level of service to all end-hosts participating in the system, including those participating via the host-specific ports on the SRP routers themselves.

Accordingly, a novel method and system is described for a standby router protocol that provides increased participation in the SRP without a concomitant increase in the amount of SRP messaging and resulting network traffic. From the foregoing description, those skilled in the art will recognize that many other variations of the present invention are possible. In particular, while the present invention has been described as being implemented in a network comprising one or more routers R3 210 and R4 220, some of the logic may be distributed in other components of a network or inter-network application.

For example, embodiments of the invention may be represented as a software product stored on a machine-accessible medium (also referred to as a machine or computer-readable medium, or a processor-readable medium). The machine-accessible medium may be any type of magnetic, optical, or electrical storage medium including a diskette, CD-ROM, memory device (volatile or non-volatile), or similar storage mechanism. The machine-accessible medium may contain various sets of instructions, code sequences, configuration information, or other data. As an example, the procedures described herein for an SRP for Domain Master VLANs 300 and an SRP for Host-Specific Ports 400 can be stored on the

machine-accessible medium. Those of ordinary skill in the art will appreciate that other instructions and operations necessary to implement the described invention may also be stored on the machine-accessible medium.

Thus, the present invention is not limited by the details described. Instead,
5 the present invention can be practiced with modifications and alterations within the spirit and scope of the appended claims.

10
15
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
100
105
110
115
120
125
130
135
140
145
150
155
160
165
170
175
180
185
190
195
200
205
210
215
220
225
230
235
240
245
250
255
260
265
270
275
280
285
290
295
300
305
310
315
320
325
330
335
340
345
350
355
360
365
370
375
380
385
390
395
400
405
410
415
420
425
430
435
440
445
450
455
460
465
470
475
480
485
490
495
500
505
510
515
520
525
530
535
540
545
550
555
560
565
570
575
580
585
590
595
600
605
610
615
620
625
630
635
640
645
650
655
660
665
670
675
680
685
690
695
700
705
710
715
720
725
730
735
740
745
750
755
760
765
770
775
780
785
790
795
800
805
810
815
820
825
830
835
840
845
850
855
860
865
870
875
880
885
890
895
900
905
910
915
920
925
930
935
940
945
950
955
960
965
970
975
980
985
990
995